

REMARKS

Claim amendments

In reviewing the claims, Applicant notes that claims 4 and 22 are inconsistent with the subject matter of claims 1 and 19. Accordingly, Applicant cancels those claims.

Scherr

Many data storage systems features a cache management system that copies some of the content from a disk into a cache memory.

In operation, when a user requests content from a server, the server checks to see if the content is already in cache memory. If so, the server serves the content directly from cache.

Scherr teaches a system along the lines of the foregoing, but in which different data storage systems use different cache management algorithms.¹ Among the algorithms *Scherr* teaches is the "time-currency method."²

In the "time-currency" algorithm, a server updates the cached web page upon lapse of a specified interval. *Scherr* teaches a particular example in which a web page containing stock prices is refreshed every fifteen minutes.³

It is apparent, however, that the time-currency method is not linked to whether or not the web page stored on the server is obsolete or not. The time-currency method blindly refreshes a web page every fifteen minutes, regardless of whether the web page has been changed in those fifteen minutes.

For example, in the context of *Scherr*'s example, if the stock's price remains unchanged for fifteen minutes, the time-currency method will upon detecting the lapse of fifteen minutes,

¹ *Scherr*, U.S. Patent No. 6,799,248, FIG. 2B and col. 3, lines 54-60.

² *Scherr*, column 6, lines 5-24.

³ *Scherr*, column 6, lines 18-24 ("A time-currency method of cache management can be configured to refresh certain pages with one frequency, say every 15 minutes, during trading hours...").

update the stored web page anyway. This has the disadvantage of wasting bandwidth refreshing a web page that is still current.

Alternatively, if, in *Scherr*, the stock price were to change one minute after an update, then the stored web page would be obsolete for fourteen minutes, i.e. until the next update. This has the disadvantage of creating a fourteen minute gap during which a server provides an obsolete web page to users.

The operation of *Scherr* is thus similar to Applicant's description of a prior art cache manager:⁴

"In such cases, the cache manager 24 need only monitor the time and transmit a request for an updated web-page at the appropriate time. This is a disadvantage because, as noted above, an object may need to be replaced prematurely, either because an unpredictable event occurred or because an event occurred at an unpredictable time. Since the conventional cache manager 24 will only replace a web page on its designated date, the premature replacement of a web page is difficult."

Section 102 rejection of claim 1

Claim 1 recites

"implementing programmable rules...each programmable rule defining a triggering event ..., the occurrence of the triggering event being indicative of an obsolete portion of said web page stored in said corresponding cache server."

As best understood, the Office regards claims 1's "triggering event" as corresponding to the lapse of the sampling interval in *Scherr*'s time-currency algorithm. Thus, in the Office's view, every time fifteen minutes goes by, a "triggering event" within the meaning of claim 1 occurs.

However, claim 1 requires not just some triggering event for refreshing a web page, but a triggering event that is "indicative of the existence of an obsolete portion" of the stored web page.

⁴ Applicant's specification, page 5, line 22-page 6, line 2.

In *Scherr*, the passage of fifteen minutes triggers the refreshing of a web page. But this triggering event is only indicative of the passage of fifteen minutes. It is not “indicative of the existence of an obsolete portion” of a web page.

As an example, suppose at $t=0$ minutes, a web page containing a stock price were refreshed. Suppose that at $t=2$ minutes, the stock price were to change. This change in stock price would result in “the existence of an obsolete portion of said web page stored in” the cache server.

Given that there now exists an obsolete portion of a web page, the question becomes what does *Scherr* do about it. The answer, of course, is that *Scherr* does nothing whatsoever. The next update of the web page would still occur at $t=15$ minutes, just as if the stock price had not changed at all. This is because the triggering event, i.e. the passage of fifteen minutes, depends *solely on the passage of time*. It is not, in any way linked to the “existence of an obsolete portion of said web page” as required by claim 1.

As another example, suppose again that at $t=0$ minutes a web page containing a stock price were refreshed. Now suppose that the stock is not an actively traded one, and that after fifteen minutes, the price still has not changed. In that case, the web page stored on the cache server would be current. There would be no “obsolete portion of said web page.”

However, in this second case, *Scherr* would still update the web page at $t=15$ minutes, *even though the web page stored on the cache server would still be current*. Thus, the triggering event, namely the passage of fifteen minutes, would have occurred even in the absence of “an obsolete portion of said web page.”

In fact, the alleged “triggering event” in *Scherr* is much like the hourly chime on a clock. It occurs relentlessly, regardless of whether or not any portion of the webpage is obsolete.

It is evident from the first example that in *Scherr*, a triggering event would *not* occur upon the change of a stock price. It is also evident from the second example that in *Scherr*, a

triggering event would occur *even if the web page on the cache were current*. Given these two observations, it does not appear possible that the triggering event, namely the passage of a selected time interval, is in any way “indicative of the existence of an obsolete portion of said web page”.

Claim 19 recites a limitation similar to claim 1 and is patentable for at least the same reason.

Claim 13 is amended to conform to claim 1. Accordingly, claim 13 is patentable for at least the same reasons as claim 1.

Claims 2-12, 14-18, and 20-28 all depend on claims 1, 13, and 19 respectively, and are therefore all patentable for at least the same reasons as those claims.

Section 103 rejection of claims 6, 24, 7, 17, and 25

In response to the §103 rejection of the independent claims, Applicant submits the attached declaration under 37 CFR 131.

As indicated in the declaration, Applicant conceived and reduced the claimed invention to practice before the earliest priority date of either *Kredo* or *Nashed*. Accordingly, Applicant requests that both *Kredo* and *Nashed* be disqualified as references for claims 6, 7, 17, 25, and 25.

Section 102 rejection of claim 3

Claim 3 recites the further limitation that “implementing said programmable rules comprises interpreting a script containing instructions for defining a rule.”

Applicant agrees that implementation of the time-currency algorithm in *Scherr* inherently requires some sort of programming instructions. However, this does not mean that the time-currency algorithm is implemented by interpreting a script. For example, one could implement the time-currency method using an executable image generated by compiled source code. This would not require interpretation.

Scherr is silent on the question of how the time-currency algorithm is implemented. In particular, nothing in *Scherr* suggests that interpreted, rather than compiled source code was used. Applicant reminds the Office that a reference is good for what it teaches, not for what it could have taught but did not.

The distinction between interpreted scripts and compiled source code is a distinction that makes a difference. An interpreted script, although less efficient to execute, is easily modified by users who may wish to modify a triggering event. In contrast, compiled code, although quickly executed, is essentially cast in stone once it is compiled. When a rule is implemented in compiled code, it is all but impossible for a user to modify the rule in any way.

As is clear from the specification, at the time of filing, Applicant already appreciated the advantage of providing a way to interpret a programmable script:

To address the foregoing disadvantages, a system 10 according to the invention provides communication between the cache manager 24 and a programmable script 36. The programmable script 36 can be a set of JavaScript instructions provided by a programmer. The script 36 can thus cause the cache manager 24 to update selected constituent objects of a web page upon the occurrence of a programmer-defined triggering event. One such triggering event can, of course, be the passage of a selected amount of time. However, because the script 36 is freely programmable, the triggering event can be any event that can be defined by the script 36. For example, if more than five clients order the same item, the script 36 can issue a request for an update from the origin server 12 even though the designated replacement time for a particular object may not have arrived. The programmable script 36 thus liberates the cache manager 24 from relying solely on the passage of time as a triggering event for replacing constituent objects of web-pages.⁵

Applicant further recognizes that in *Scherr*, one can specify a different triggering event for each of the cache management systems.⁶ However, this does not involve interpretation of a programmable script. This is merely interacting with a program to change a parameter.

⁵ Applicant's specification, page 6, lines 3-15.

⁶ *Scherr*, FIG. 2b.

Section 102 rejection of claim 5

Claim 5 recites the additional limitation of detecting a triggering event by “detecting the receipt of an updated portion of said web page.”

The cited text lists several types of what could be regarded as triggering events:

Depending on the configuration(s) selected, the system may manage data or subsets of data in a storage cache on the basis of time-currency, page usage frequency, charging considerations, pre-fetching algorithms, data usage patterns, store-through methods for updated pages, least recently used method, B-tree algorithms, or indexing techniques including named element ordering, among others.⁷

None of these putative “triggering events” involve “detecting the receipt of an updated portion of said web page.”

Section 102 rejection of claim 8

Claim 8 recites the additional limitation of “establishing communication with an origin server and causing said particular cache server to request said update therefrom.”

The Office suggests that *Scherr* teaches the foregoing claim limitation in FIG. 2a:

“If, in FIG. 2a, at decision block 30 it is determined that the data is not already in the cache...a request will be made to fetch the data from the network at step 34.”

However, this text merely teaches obtaining data from the network. There is no indication that obtaining data from a network involves requesting an update from an origin server.

Nor is an origin server inherent in any act of retrieving data over a network. For example, many personal computers are connected to network drives that are accessed just like local drives. In such cases, obtaining data from a network amounts to issuing a copy command, and does not involve an origin server at all.

⁷ *Scherr*, col. 3, lines 54-60.

Section 102 rejection of claims 9 and 10

Claims 9 and 10 recite limitations on the location of the cache memory relative to the origin server.

The Office suggests that column 5 of *Scherr* describes locations of cache memory relative to the origin server.

Applicant is unable to identify any text in column 5 that teaches the location of cache memory relative to an origin server. The cited passage does describe the possibility that the cache memory is a RAID system. However, it does not describe the location of the RAID system relative to anything that could be construed as an origin server.

Applicant requests that the Examiner quote verbatim the text from column 5 that is believed to teach the limitations of claims 9 and 10.

SUMMARY

Now pending in this application are claims 1-3, 5-21, and 23-28. Of these, claims 1, 13, and 19 are independent. Please apply the \$525 charge for the extension of time fee along with any other charges or credits to deposit account 06-1050, referencing Attorney Docket No. 11125-014001.

Respectfully submitted,

Date: October 1, 2007



Faustino A. Lichauco
Reg. No. 41,942

Fish & Richardson P.C.
225 Franklin Street
Boston, MA 02110
Telephone: (617) 542-5070
Facsimile: (617) 542-8906